

Attorney Docket No.: 82771.P259

Patent

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

TITLE OF THE INVENTION

METHOD FOR IMPROVING SIGNALING EFFICIENCY
AND PERFORMING SERVICE LOAD BALANCING
IN A CONNECTION ORIENTED NETWORK

INVENTORS

JAMES V. LUCIANI

Prepared by

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026
(408) 720-8598

"Express Mail" mailing label number EL034432708 US

Date of Deposit OCTOBER 29, 1999

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231

Mary E. Wainner
(Typed or printed name of person mailing paper or fee)

Mary E. Wainner
(Signature of person mailing papers or fee)

COPYRIGHT NOTICE

Contained herein is material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent disclosure by any person as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all rights to the copyright whatsoever.

BACKGROUND OF THE INVENTION

Field of the Invention

10 The present invention is related to data communications. In particular, the present invention is related to a method for resolving a request for information in a multiprotocol internetwork environment operating over Non-Broadcast Multiple Access (NBMA) subnetworks.

Description of the Related Art

15 The Next Hop Resolution Protocol (NHRP) is described in Luciani, J., Katz, D., Piscitello, D., Cole, B., and Doraswamy, N., "NBMA Next Hop Resolution Protocol (NHRP)", IETF RFC 2332, April, 1998. With reference to Fig. 1, the NHRP allows a source station 110 (a node, host, switch, router, etc.) in
20 a multiprotocol internetwork 100 to communicate with a destination station 130 over a Non-Broadcast, Multiple Access (NBMA) subnetwork 105. The NHRP provides source station 110 with the capability to determine the NBMA address of

the NBMA next hop toward the destination station, through the exchange of NHRP resolution requests and resolution replies. As pointed out in RFC 2332, the NBMA subnetwork may be non-broadcast because it does not support broadcasting (e.g., an X.25 or ATM subnetwork) or broadcasting is not possible, for example, in the case of a Local Area Network (LAN) that supports a large number of stations. In Fig. 1, the NBMA next hop toward the destination station is router 120 since it is the closest router to destination station 130 and provides egress from the NBMA subnetwork. It should be noted that terms used herein to describe the present invention, such as internetwork layer, server, client, and station, are to be interpreted in a manner consistent with the definition and use of such terms as provided in RFC 2332.

In accordance with the NHRP, Next Hop Servers (NHSs) are provided in the NBMA subnetwork which are capable of responding to NHRP resolution requests. In Fig. 1, egress router 125 also functions as a NHS, and serves one or more destination stations, such as destination station 130. Likewise, ingress router 115 must function as a NHS for station 110.

An NHS builds and maintains a data structure that contains internetwork layer address (e.g., an Internet Protocol address) to NBMA subnetwork layer address resolution information. The table may be built and managed in accordance with techniques known to those of ordinary skill in the related arts. For example, a station may send a NHRP registration request to a NHS serving the station. The NHRP registration request contains internetwork layer address to

NBMA subnetwork layer address resolution information that is then stored in the table maintained by the NHS.

A station that is a client of the NHRP service is known as a NHRP Client, or simply NHC. The NHS with which a NHC communicates to provide NBMA
5 next hop information is the serving NHS for the NHC. In Fig. 1, NHS 115 serves station (NHC) 110, and NHS 125 serves station (NHC) 130 in most cases. For a serving NHS to supply address resolution information to a NHC, a continuous link of NHSs must exist along a path in the NBMA subnetwork between the NHC making the NHRP resolution request and the destination NHC, e.g., NHSs 115,
10 120 and 125. In accordance with RFC 2332, the last NHS along the path within the NBMA subnetwork is the serving NHS. That is, NHRP resolution requests are not forwarded to destination station/NHCs but are processed by the serving NHS. However, each NHC also maintains a table of internetwork layer address to NBMA address resolution information that it obtains from NHRP resolution
15 replies, manual configuration, or through mechanisms outside the scope of the NHRP. Destination NHCs may be constrained on resource (e.g., SAR VPI/VCI), and there is no existing mechanism to communicate that fact to the serving NHS which would normally reply to a resolution request on behalf of the NHC. Thus, the only way that a source station would find that out would be to attempt a
20 connection setup and fail which is time consuming and resource intensive. By allowing the destination NHC to reply for itself, since it is in the best situation to know whether it has resource enough for the connection, signaling efficiency is

gained since no additional connection setup attempt will be tried by the source station. Moreover, in the situation where multiple NHCs register with a given serving NHS for the same set of NBMA subnetwork addresses (e.g., for the same set of ATM attached servers/services), the NHS may perform service load

- 5 balancing by forwarding the resolution request to a particular NHC and if that NHC then NAKs the resolution request, the serving NHS may offer the request to another NHC which is not currently busy. Further the NHS may choose to offer the request in some scheduled fashion (e.g., round ^{robinning} ~~robin~~) to each of the appropriate NHCs in turn.

BRIEF SUMMARY OF THE SEVERAL VIEWS OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the following figures, in which:

5

Figure 1 is a prior art diagram of a multiprotocol internetwork environment operating over an NBMA subnetwork.

Figure 2 illustrates an embodiment of the present invention.

10

Figure 3 illustrates a multiprotocol internetwork environment in which an embodiment of the present invention may be utilized.

DETAILED DESCRIPTION OF THE INVENTION

Described is a method providing for a NHC to respond to a NHRP resolution request. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It
5 will be apparent, however, to one of ordinary skill in the art that the present invention may be practiced without these specific details. In other instances, well-known architectures, steps, and techniques have not been shown to avoid unnecessarily obscuring the present invention.

With reference to Fig. 2, an embodiment of the present invention 200
10 proceeds as follows. The process starts as 205, wherein source station 110, given the internetwork layer address of destination station 130, seeks to resolve the NBMA subnetwork address of a path to station 130. At 210, if the NBMA subnetwork address information for the path to station 130 is already available in the address resolution table maintained at the source station, then that information
15 is utilized by the source station in communicating with station 130 and the resolution process is done at 215. Otherwise, the source station creates at 220 a NHRP resolution request packet comprising the internetwork layer address of the destination station as the destination address, the internetwork layer address of the source station as the source address, and the NBMA subnetwork address
20 information for the source station. Source station 110 then transmits the NHRP resolution request packet at 225 to the nearest NHS (i.e., ingress router 115) along

the routed path of NHSs within the NBMA subnetwork 105, toward the destination station 130.

The NHS 115 receives the NHRP resolution request at 230, examines at 235 whether it serves destination station 130, and if not, forwards the resolution request at 225 to the next NHS in the routed path to destination station 130, in this case, intermediate NHS 120.

The process continues in this manner until NHS 125 receives and examines the destination internetwork address in the resolution request packet and determines at 235 that it serves destination station 130. NHS 125 then considers at 240 whether it should respond to the NHRP resolution request on behalf of its NHS, destination station 130, or whether it should forward the resolution request on to the destination station so that the destination station may respond directly to the resolution request. If the serving NHS determines it should reply to the resolution request on behalf of its NHC, destination station 130, then the serving NHS formulates, and transmits at 260, a positive NHRP resolution reply that resolves the NBMA address information for destination station 130 on the destination station's behalf. The NHRP resolution reply packet contains the address resolution information for the destination station and is sent back to the source station. It should be noted that if the destination station is not on the NBMA subnetwork, as is the case in multiprotocol internetwork 100, the next hop subnetwork layer address will be that of the egress router 125 through which packets addressed to the destination station are forwarded.

In determining at 240 whether the serving NHS or the NHC/destination station should respond to a NHRP resolution request, in one embodiment of the present invention, a destination station may inform its serving NHS that it wishes to directly reply to NHRP resolution requests received at the serving NHS from another station. One manner in which the NHC may inform the serving NHS of its wish to directly reply to NHRP resolution requests is to indicate via a unique time/length/value (TLV) extension to a NHRP registration request packet that it wishes to receive and respond to any NHRP resolution requests that the serving NHS receives, rather than allowing the serving NHS to respond to the resolution request on behalf of the NHC/destination station.

In another embodiment, different TLV values may be specified in a NHRP registration request sent from the NHC to the NHS to indicate certain situations under which it is applicable for or desired that the NHS forward NHRP resolution requests to the NHC. For example, the NHC may specify via the TLV extension part of a NHRP registration request packet certain resource constraints on the NHS during which time the NHS is to forward NHRP resolution requests about the NHC to the NHC rather than responding to the request on behalf of the NHC. Alternatively, a NHC may choose whether to reject or accept a NHRP resolution request forwarded to it from its serving NHS based on an agreed upon configuration, e.g., through manual (user) configuration, management signaling, etc, or via information obtained by another protocol, for example, via a routing

protocol. Under these approaches, communication of such configuration via the TLV extension in a NHRP registration request packet is not necessary.

With respect to Fig. 3, in yet another embodiment of the present invention, when multiple NHCs, e.g., stations 130, 140 and 150, transmit NHRP registration requests to the serving NHS 125 to register for the same destination address or set of destination addresses (e.g., for the same set of ATM attached servers/services), the NHS may perform load balancing by directing a NHRP resolution request to a certain one or more of the NHCs. The choice of which one or more of the NHCs to offer the request to and in what manner is a local matter. An appropriate NHC may be selected according to any one of a number of means. For example, a round robin, weighted round robin, or some other arbitration scheme may be employed to select the appropriate NHC to receive the NHRP resolution request. Moreover, the NHC selected may decide to accept or reject the NHRP resolution request based on certain criteria, such as resource availability, source address of the source station, or any other well known packet filtering technique. This approach often is advantageous given that a particular NHC generally is in a better position to determine if it is willing to receive communications from a source station than its serving NHS.

In the embodiment illustrated in Fig. 3, if a given NHC has indicated to its serving NHS that ^{it} is capable of receiving and responding directly to a NHRP resolution request, but nevertheless decides to reject a NHRP resolution request, then the serving NHS, upon receiving a negative NHRP resolution reply from the

NHC, may elect to forward the NHRP resolution request on to another NHC that satisfies the request criteria. In this way, an improved load balancing capability and better set up connection signaling efficiency is achieved.

CONFIDENTIAL